STAP Training through Knowledge-Aided Predictive Modeling

Nathan A. Goodman and Prashanth R. Gurram
Department of Electrical and Computer Engineering, The University of Arizona
1230 E. Speedway Blvd., Tucson, AZ 85721
goodman@ece.arizona.edu

Abstract – In this paper, we investigate a spectral-domain approach to estimating the interference covariance matrix used in space-time adaptive processing. Traditionally, an estimate of the interference covariance matrix is obtained by averaging the space-time covariance matrices of multiple range bins. Unfortunately, the spectral content of these data snapshots usually varies, which corrupts the covariance estimate for the desired range. We propose to use knowledge sources to identify angle-Doppler spectral regions having the same underlying scattering statistics. Then, we use real-time data to form a synthetic aperture radar image, which is inherently an estimate of non-moving ground clutter. We then average the SAR pixels within each homogeneous region. The resulting clutter power map is used, along with knowledge of the radar system and scenario geometry, to compute the interference covariance matrix. Using simulated data, we demonstrate the potential performance of such a technique, demonstrate its dependence on accurate space-time steering vectors, and provide an example of using data to compensate for imperfect knowledge.

I. Introduction

Conventional space-time adaptive processing (STAP) training requires significant sample support for acceptable performance [1-3]. When training data from nearby range bins are independent and identically distributed, a good estimate of the interference covariance matrix can be obtained by averaging covariance matrices formed from data snapshots taken from multiple range bins. Reduced-rank methods reduce the sample support necessary for covariance training, but can still suffer from poor covariance estimation in non-stationary clutter and dense target environments [4-5].

In this paper, we present preliminary analysis on an alternative training method using known properties of the scattering background. In this method, we propose to use knowledge of the clutter scene to perform a priori identification of regions over which the scattering statistics are homogeneous. Then, we propose to use the boundaries of these homogeneous regions along with real-time synthetic aperture radar (SAR) imagery to estimate clutter's range/cross-range power profile. Due to the speckle effect, SAR pixels cannot be used directly as an estimate of clutter. Instead, speckle is removed by averaging all SAR pixels within each of the regions identified as being homogeneous. Then, the averaged clutter profile is transformed to the clutter covariance matrix based on predictable space-time steering vectors. Finally, we use real-time data to compensate for errors in the clutter covariance calculation due to finite-precision knowledge of the radar system parameters and flight geometry.

In this paper, we present a description of our proposed technique and initial results on some of the hurdles that must be overcome in order to achieve good performance. In Section II, we describe our underlying algorithm in more detail. In Section III, we present our simulated results. We show that in the limiting case of perfect knowledge, our approach produces near-ideal performance. We also list some specific assumptions that are inherent in the claim of perfect knowledge. We present analysis on the performance degradation that is observed when the knowledge used to compute the space-time steering vector for each clutter patch is imperfect. Specifically, we analyze the situation where we have imperfect knowledge of the radar platform's crab angle. We then present an example of how data can be used to compensate for these unknowns. This example involves using superresolution to estimate the space-time clutter ridge, which is then used to estimate crab angle. The estimated crab angle is then used to compute improved steering vectors that nearly restore the perfect-knowledge performance. In Section IV, we discuss some of the future research directions that we envision for this technique, and in Section V we make our conclusions.

II. SAR-Based Spectral Averaging Approach

Currently, STAP training always occurs directly in the measurement domain [1-2]. Multiple space-time data *snapshots* are taken from range bins other than, but usually nearby to, the current range under test (RUT). These snapshots are then used to form the sample covariance estimate of the interference covariance matrix according to [1]

$$\hat{\mathbf{R}}_I = \frac{1}{N_s}\sum_{i=1}^{N_s}\mathbf{d}_i\mathbf{d}_i^H \tag{1}$$

where $\mathbf{d}_i$ is the i^{th} secondary data snapshot, N_s is the number of secondary data snapshots, and the superscript H denotes conjugate transpose. Once an estimate of the interference covariance is obtained, the (estimated) optimum weight vector for detecting a target is well known as

0-7803-8234-X/04/$17.00

$$\hat{\mathbf{w}}_o = \kappa \hat{\mathbf{R}}_I^{-1} \mathbf{s} \tag{2}$$

where κ is an arbitrary constant and $\mathbf{s}$ is the space-time steering vector for a moving target with a given radial velocity and along-track position.

If the number of pulses in a coherent processing interval (CPI) is M, and the number of spatial channels is N, then the number of space-time measurements is MN, and the dimension of the covariance matrix estimate in (1) is MN by MN. If $2MN$ independent and identically distributed data snapshots are available for performing the estimate in (1), then performance approaches the ideal, known-covariance case to within about 3 dB [3]. Unfortunately, this performance level can rarely be achieved in practice due to non-stationarity of the training data. In reality, the along-track scattering profile of ground clutter varies with range, as does the received power due to its dependence on the fourth power of range, R^4. Additional targets within the training region also often corrupt the covariance estimate [6], and aircraft crab [1,6] and non-linear arrays [7] cause the space-time clutter ridge to vary with range. Some discussions of real-world effects and their impacts on STAP performance can be found in [2,6].

Although the above *sample matrix inversion* (SMI) technique can suffer from degraded performance due to non-stationarity of the training data, there are some benefits to the approach that should be mentioned. Because the interference statistics are estimated from real-time data, any jammers that are present will also be estimated in the training process. Hence, the SMI approach inherently estimates the statistics of clutter, additive white gaussian noise, and jammers in one well-defined process. The SMI approach also inherently includes channel mismatches, or other system characteristics, in the covariance estimation, although mismatches will still be present in the target steering vector, s.

An alternative method of computing the interference is possible if the clutter RCS background and platform kinematics are known. This approach is usually used to calculate the true covariance matrix for performance analysis in simulated scenarios. If a data sample collected by the radar system at some time, t, and along-track spatial position, r_x, due strictly to the return from ground clutter is $d_c(t,r_x)$, then the covariance of two measurements can be represented as

$$\begin{aligned} R_c(\tau,\chi) &= E\left[d_c(t,r_x)d_c^*(t+\tau,r_x+\chi)\right] \\ &= \int_X \sigma_c(x)h(t,r_x,x)h^*(t+\tau,r_x+\chi,x)\,dx \end{aligned} \tag{3}$$

where τ and χ are the delay and spatial separation, respectively, between the two measurements, $\sigma_c(x)$ is the along-track clutter RCS profile within a range bin, $h(t,r_x,x)$ is a function describing the normalized measurement obtained at time, t, and spatial position, r_x, due to a scatterer located at x (azimuth), and the integration is performed over the azimuth illumination width of the RUT. When the integration in (3) is approximated with a summation, and the received signal is sampled in space and time, the matrix representing all space-time lags is the clutter covariance matrix presented in [1]:

$$\mathbf{R}_c = \sum_{k=1}^{N_c} \sigma_k \mathbf{v}_k \mathbf{v}_k^H \tag{4}$$

where N_c is the number of clutter *patches*, $\mathbf{v}_k$ is the space-time steering vector to the k^{th} patch, and σ_k is the received per-element, per-pulse power for the k^{th} patch. Although it is impossible in practice to obtain perfect knowledge of the σ_k's in (4), a priori knowledge could be used to help obtain an estimate of the observed clutter scene. This estimate could then be used in conjunction with accurate knowledge of platform kinematics to compute an estimate of the clutter covariance matrix.

We first consider estimation of the clutter power profile (the σ_k's) in (4). To begin, we point out that SAR inherently provides an estimate of ground clutter. SAR creates a map of stationary-scatterer reflectivity versus range and azimuth. Unfortunately, the values in a SAR image cannot be used in (4) directly due to the speckle phenomenon – each pixel in a SAR map represents the sum of many complex scatterers within a resolution cell; hence, a single SAR pixel likely does not represent its true average scattering value. However, if multiple SAR pixels from the same surface type are available, they can be averaged to get an accurate, representative value of that surface's RCS. For example, the boundaries of a grassland, lake, or agricultural area could be identified, and we would expect the scattering within the boundaries of one of these features to adhere to the same probability distribution. If the boundaries of such homogeneous regions, or segments, can be identified, then all pixels within the boundaries can be averaged. Then, each SAR pixel can be replaced with its region's average power value for use as a clutter patch power coefficient in (4).

Reflections from ground clutter produce a random signal at the radar receiver. The signal is random because, although the average power reflected from a clutter patch may be known, the speckle phenomenon causes the scattering from any single patch to be random. Hence, the average RCS profile is really the power spectral density of the clutter random process. Furthermore, just as the autocorrelation function and power spectral density are equally valid representations of the same random process in the time and frequency domains, respectively, the clutter covariance matrix and range/cross-range RCS profile are equally valid representations of clutter in the data and spectral domains. However, both representations are usually not known, and the data-domain representation is more easily computed. In general, it does not make sense to try to estimate the random process in the spectral domain because we need average power but it is not clear which spectral components it is appropriate to average together. Hence, we arrive at our prior conclusion that it is not possible to directly use a SAR image as the clutter power spectral density.

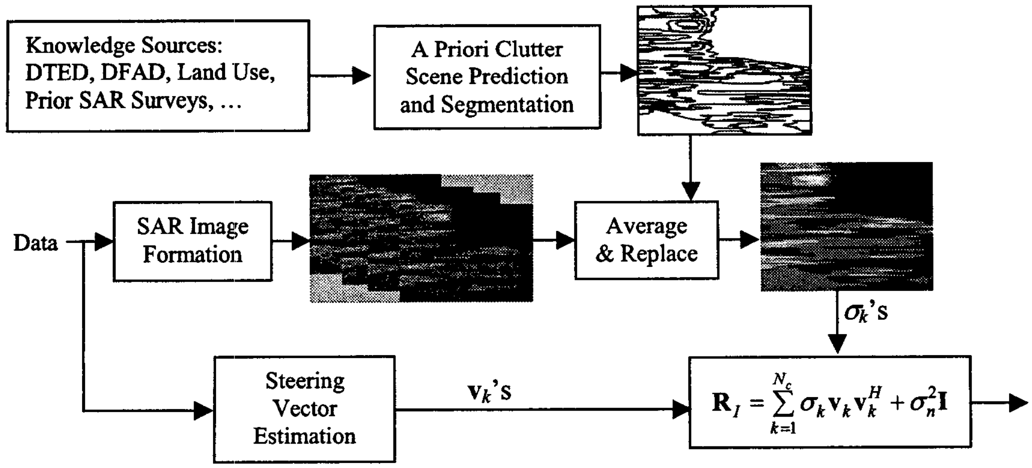

Fig. 1. Block diagram of the proposed, spectral-averaging-based processing architecture.

Our proposed technique, however, does in fact perform averaging of the clutter random process in the spectral domain rather than in the measurement domain as in conventional STAP. Pixels that adhere to the same probability distribution are averaged to reduce speckle, but as just mentioned above, we must identify regions over which averaging is appropriate. We suggest that this is an ideal use of knowledge-aided radar signal processing. We suggest that a priori knowledge from many sources – previous SAR surveys, DTED, land cover, and other known features – could be used to predict homogeneous spectral regions prior to SAR image formation. Then, these segmentations could be used to average a real-time SAR image and arrive at a spectral-domain representation of clutter.

Of course, in order to transform between the spectral and data domains, we see in (4) that space-time steering vectors are needed. Ideally, these can be computed based on available ownship data. The temporal steering vectors for each clutter patch (SAR pixel) are known since they are used to form the SAR images on each channel. The spatial steering vectors can also be computed, but our initial experiments have shown that performance is very sensitive to even small errors in ownship knowledge. For example, the clutter notch formed by the optimum space-time processor is very narrow, and imperfect knowledge of array misalignment, or crab, can cause the clutter notch to miss the actual clutter ridge. Therefore, we conclude that real-time data must be used to adaptively correct for imperfect knowledge of crab angle, channel imbalance, or other factors.

Figure 1 shows a block diagram of the proposed processing approach. Knowledge sources are used to predict the observed clutter scene for a known flight profile. Then, this prediction is used to divide the observed scene into homogeneous scattering regions. Data collected by the radar system are used to form SAR images, which are then segmented according to the a priori predictions that have been adjusted to compensate for any deviations in the flight scenario. After the pixels within each segment have been averaged, the pixel values can be used as the power coefficients in the covariance matrix calculation. The data are also used to estimate space-time steering vectors. For example, as we will demonstrate in the next section, real-time data can be used to obtain an accurate estimate of the radar platform's crab angle, which can then be used in the steering vector computation. After estimation of the clutter power profile and steering vectors, the clutter covariance matrix is calculated and a known noise floor is added.

III. Simulations

In this section, we demonstrate the potential performance of the technique described above. We show that under ideal conditions, performance is significantly better than the SMI method, but in order to achieve good performance in practice, very accurate knowledge of the space-time steering vectors for the clutter patches is required. This, in turn, places strict requirements on the accuracy of ownship and scenario knowledge. The required accuracies probably cannot be met without using real-time data, so we also demonstrate an example of using real-time data to estimate platform crab angle.

Figure 2 demonstrates that the spectral-averaging-based approach provides near-optimal performance in the limiting case of perfect knowledge. A STAP simulation was performed with parameters similar to those given in [7]. The data cube, calculated using the clutter patch model of [1], consisted of noisy space-time data for $2MN$ range bins plus the RUT and a few guard cells. The aircraft crab angle was 3.5 degrees. Figure 2 shows SINR loss for the known-covariance case, SMI, and our proposed approach. The heterogeneity of the clutter scene and the range-dependent clutter locus reduce SMI performance below the 3-dB loss that would be expected for $2MN$ snapshots of independent and identically distributed data. The proposed approach, however, achieves performance

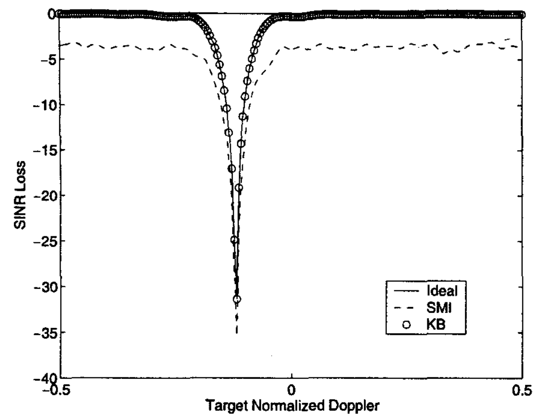

Fig.2. SINR Loss for the ideal, SMI, and proposed (KB) covariance estimates under perfect scenario knowledge.

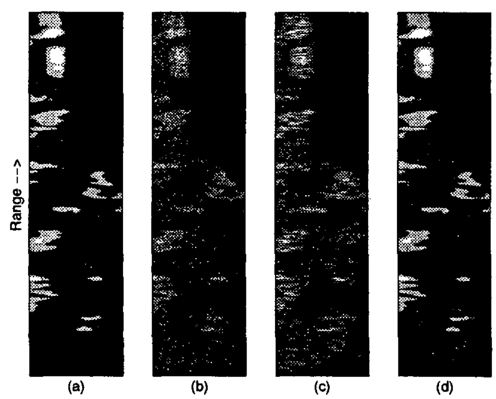

Fig. 3. The (a) ideal, (b) speckled, (c) SAR-reconstructed, and (d) averaged clutter power profiles.

that is indistinguishable from the ideal, known-covariance performance.

In order to achieve the performance shown in Fig. 2, perfect knowledge of the radar system and scenario were assumed. In particular, we assumed that the noise level and gain of each channel were known constants, that aircraft crab angle was known and constant throughout the CPI, and that the scattering coefficients of the clutter patches were circularly complex Gaussian random variables with known variance. Since the underlying variance of each clutter patch was known, homogeneous scattering segments could be identified perfectly, resulting in optimal averaging performance. Furthermore, there were no moving targets in the data. All other system parameters, such as altitude or pulse repetition frequency, were also known.

Figure 3 shows the range/cross-range scattering profile that was used. In Fig. 3a, the ideal clutter RCS profile is shown. In Fig. 3b, we show the amplitude of the scattering coefficients after the speckle phenomenon was added. In Fig. 3c, we see the SAR image formed by the first spatial channel, and in Fig. 3d, we see the resulting power profile after the known homogeneous regions are averaged.

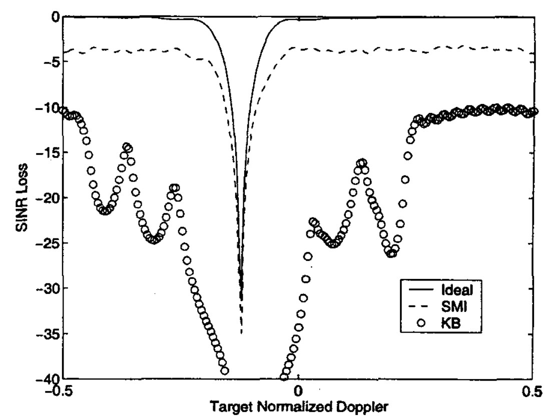

Fig. 4. SINR Loss for the ideal, SMI, and proposed (KB) covariance estimates under incorrect platform crab knowledge.

In practice, the available ownship and scattering knowledge will have limited accuracy. For example, the platform's crab angle will be known approximately, but not perfectly. Furthermore, we have found that our approach is sensitive to imperfect knowledge assumptions. Figure 4 demonstrates such sensitivity. In producing Fig. 4, the data cube was generated for a platform with an actual crab angle of 3.5 degrees. In computing the space-time steering vectors necessary for transforming the clutter power coefficients to the clutter covariance matrix, it was assumed that the platform's INS reported a crab angle of 3.3 degrees. As can be seen from Fig. 4, this slight error drastically reduces performance. The reason for reduced performance is that the clutter notch that is formed by the optimum filter is very narrow (the clutter ridge is infinitely narrow since we have not yet modeled any intrinsic clutter motion). Even small errors can cause the clutter notch to miss the clutter ridge entirely; hence, the clutter is severely undernulled.

Figure 5 shows the eigenvalues of the ideal, SMI, and knowledge-aided interference covariance matrices. It is seen that the knowledge-aided approach represents the true eigenvalues of the clutter-plus-noise subspace more accurately than SMI despite the fact that SMI shows better performance in Fig. 4. This is further evidence that the poor performance of the knowledge-aided approach in Fig. 4 is due to inaccurate transformation from the spectral domain to the data domain. In other words, the steering vectors are the root of the problem, not the clutter patch power coefficients.

Figure 4 demonstrates that the spectral-averaging-based approach will not be effective without using real-time data to estimate practical factors that affect STAP performance. While SAR processing should produce a reasonably effective clutter power profile, performance will only be as good as the space-time steering vectors that transform that power profile to a covariance matrix.

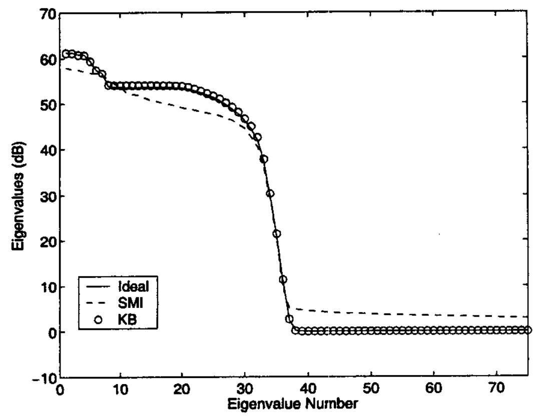

Fig. 5. Eigenvalues of the ideal, SMI, and proposed (KB) covariance matrix estimates under incorrect platform crab knowledge.

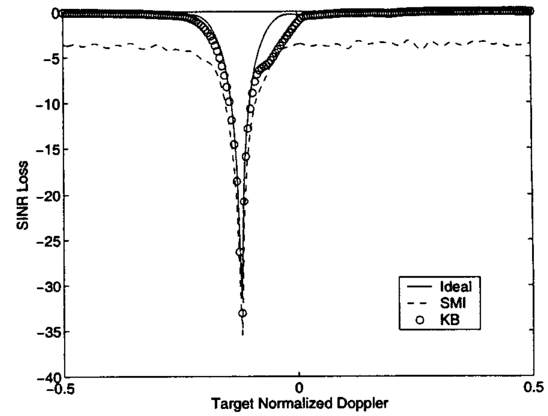

Fig. 6. SINR Loss for the ideal, SMI, and proposed (KB) covariance estimates. The KB covariance matrix was corrected using the superresolution estimate of crab angle.

In the following, we present an example of how real-time data might be used to estimate system characteristics that affect STAP performance, and how those estimates could be used to restore the performance of the proposed approach. In the next simulation, the true crab angle for the platform was again 3.5 degrees. However, in this example we used the traditional sample covariance matrix to locate the clutter ridge. Then, we computed the crab angle that would be required to produce the estimated clutter ridge and used the estimated crab angle in calculating the space-time steering vectors.

Unfortunately, a simple Fourier-based estimate of the clutter ridge is not sufficiently accurate. This approach produces a crab estimate with accuracy on the order of the angle subtended by a single resolution cell, but we require accuracy that is on the order of the width of the clutter notch. In order to achieve this accuracy, we have used a 2D superresolution approach to locate the clutter ridge in angle and Doppler.

The crab estimation process is as follows. First, we compute the traditional covariance matrix estimate as in (1). Then, we use a finite-beamwidth version of Brennan's rule [1] to estimate the dimension of the clutter subspace, known as clutter rank, r_c. We perform an eigen-decomposition of the the sample covariance matrix, and divide the eigenvectors into the clutter and white-noise subspaces:

$$\begin{aligned}\mathbf{U} &= \begin{bmatrix}\mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_{r_c} & \mathbf{u}_{r_c+1} & \cdots & \mathbf{u}_{MN}\end{bmatrix} \\ &= \begin{bmatrix}\mathbf{U}_c & \mathbf{U}_n\end{bmatrix}\end{aligned}. \qquad (5)$$

The covariance matrix corresponding to the white-noise-only subspace is

$$\mathbf{R}_n = \mathbf{U}_n \mathbf{U}_n^H, \qquad (6)$$

and the superresolution estimate of the clutter spectrum at Doppler frequency, f_d, and spatial frequency, f_x, is

$$S_c(f_d, f_x) = \frac{1}{\mathbf{v}(f_d, f_x)^H \mathbf{R}_n \mathbf{v}(f_d, f_x)} \qquad (7)$$

where $\mathbf{v}(f_d, f_x)$ is the steering vector for the given angle-Doppler frequency component. Once the angle-Doppler frequency spectrum is obtained, we select the brightest point in the spectrum, note its spatial and Doppler frequencies, and compute the crab angle. In the future, we intend to investigate other methods for computing crab angle from the clutter spectrum, such as averaging the crab estimates from multiple points or performing a least-squares fit to an elliptical clutter locus.

An interesting point to make is that, in this part of the problem, strong clutter returns are actually beneficial. For the purposes of estimating crab angle, clutter acts as the desired signal; hence, increased clutter power improves the accuracy of the clutter ridge estimate. In the simulations in this paper, the clutter-to-noise ratio (CNR) for the RUT is approximately 45 dB, with some variability in other range bins due to changing clutter background. With this CNR and using the above approach, we have been estimating crab angle with accuracies on the order of a few thousandths of a degree. With this accuracy, we can restore performance as shown in Figure 6.

IV. Comments and Future Work

We have demonstrated that in order to transform a SAR-based estimate of the clutter power spectrum into an accurate estimate of the clutter covariance matrix, the space-time steering vectors used to make the transformation must be highly accurate. The above example demonstrates how data might be used to obtain an accurate estimate of platform crab angle. There are, however, many other characteristics that will need to be estimated in a real implementation.

Another practical aspect that must be estimated is channel mismatch [2]. Each spatial channel consists of its own receiving chain. Therefore, although careful calibration is usually performed, each channel usually has a slightly different gain (amplitude and phase) and noise level. According to [2], the noise level can usually be estimated quite accurately, and differences can be modeled as differing channel gains. Although channel mismatch is not always a factor in SMI performance, we believe our approach will be more sensitive to imperfect calibration; hence, we will be investigating this factor.

Other sources of imperfect knowledge include channel mismatch due to imprecisely placed antenna elements, moving targets in the SAR image used to estimate the clutter power spectrum, and variable platform speed. Each of these factors must be accounted for in our knowledge-aided approach. Note, however, that several techniques currently being developed to mitigate the problem of moving targets in the training data can also be applied to our approach. In fact, techniques such as in [6], where range cells having roads (which are then more likely to contain targets) are removed from the training data, are simple to implement in our approach. Cells with known clutter discretes, or that are likely to contain moving targets, can simply be ignored in the spectral averaging process.

Intrinsic clutter motion (ICM) is another factor that will affect the performance of our technique. However, at this time it is not clear that ICM will have a negative effect on our technique compared to the SMI technique. We argued above that our approach is sensitive to errors in the space-time steering vectors because we were attempting to place a very narrow notch on a very narrow clutter ridge. ICM effectively increases the width of the clutter ridge, requiring a wider null to be placed in the filtering process. A wider null can be achieved simply by using a covariance matrix taper [8-9]; hence, accounting for ICM in our algorithm should be straightforward. Moreover, if the clutter ridge has finite width, then small misalignments between the clutter ridge and clutter notch will cause proportionately smaller degradation. Hence, errors in the space-time steering vectors caused by small phase errors due to channel mismatch or crab angle should have reduced impact.

We also note that our current architecture has no method for estimating and rejecting jammers. Future work must include a method for using data to estimate jamming parameters and for incorporating those estimates into the interference covariance matrix. Finally, a more fundamental framework for integrating the knowledge-aided covariance with data-based estimates is desired. This framework would help to incorporate all the possible compensations that need to be performed, would make jammer estimation an integral part of the processing architecture, and would make future improvements more feasible.

V. Conclusions

In this paper, we have described a process for estimating the clutter power spectrum from a priori identification of homogeneous scattering regions and from real-time SAR imagery. The estimated clutter spectrum is then transformed to a clutter covariance matrix using the space-time steering vector for each SAR pixel. We have also acknowledged that this approach is very sensitive to errors in the steering vectors, especially when ICM is small or nonexistent such that the clutter ridge is very narrow. In this case, small errors cause severe clutter undernulling.

We have also proposed that data-based compensations could be performed to account for finite-precision knowledge that leads to steering vector errors. The combination of data-based adaptivity with knowledge-aided estimation of the clutter spectrum holds potential for excellent performance.

References

[1] J. Ward, "Space-time adaptive processing for airborne radar," MIT Lincoln Laboratory, Lexington, MA, Tech. Rep. 1015, Dec. 1994.

[2] J.R. Guerci, *Space-Time Adaptive Processing for Radar.* Boston, MA: Artech House, 2003.

[3] I.S. Reed, J.D. Mallett, and L.E. Brennan, "Rapid covergence rate in adaptive arrays," *IEEE Trans. AES*, vol. 10, no. 6, pp. 853-863, Nov. 1974.

[4] H. Wang and L. Cai, "On adaptive spatial-temporal processing for airborne surveillance radar systems," *IEEE Trans. AES*, vol. 30, no. 3, pp. 660-670, July 1994.

[5] G.K. Borsari, "Mitigating effects on STAP processing caused by an inclined array," in *Proc. 1998 IEEE National Radar Conf.*, Dallas, TX, pp. 135-140, 1998.

[6] D.J. Zywicki, W.L. Melvin, G.A. Showman, and J.R. Guerci, "STAP performance in site-specific clutter environments," in *Proc. 2003 IEEE Aerospace Conf.*, vol. 4, pp. 2005-2020, March 8-15, 2003.

[7] M. Zatman, "Circular Array STAP," *IEEE Trans. AES*, vol. 36, no. 2, pp. 510-417, April 2000.

[8] J.R. Guerci, "Theory and application of covariance matrix tapers for robust adaptive beamforming," *IEEE Trans. Sig. Proc.*, vol. 47, no. 4, pp. 977-985, April 1999.

[9] J.R. Guerci and J.S. Bergin, "Principal components, covariance matrix tapers, and the subspace leakage problem, *IEEE Trans. AES*, vol. 38, no. 1, pp. 152-162, Jan. 2002.